

Bounds on Ramsey Games via Alterations

He Guo* and Lutz Warnke†

January 13, 2022; revised January 24, 2023

Abstract

We present a refinement of the classical alteration method for constructing H -free graphs: for suitable edge-probabilities p , we show that removing all edges in H -copies of the binomial random graph $G_{n,p}$ does not significantly change the independence number. This differs from earlier alteration approaches of Erdős and Krivelevich, who obtained similar guarantees by removing one edge from each H -copy (instead of all of them). We demonstrate the usefulness of our refined alteration method via two applications to online graph Ramsey games, where it enables easier analysis.

Keywords: Alteration method, Ramsey theory, Online Ramsey games, Deletion method

1 Introduction

The probabilistic method is a widely-used tool in discrete mathematics. Many of its powerful approaches have been developed in the pursuit of understanding the graph Ramsey number $R(H, k)$, which is defined as the minimum number n so that any n -vertex graph contains either a copy of H or an independent set of size k . For example, in 1947 Erdős pioneered the *random coloring approach* to obtain the lower bound $R(K_k, k) = \Omega(k2^{k/2})$ as $k \rightarrow \infty$, where K_k denotes a k -vertex clique, and in 1961 he developed the *alteration method* in order to obtain $R(K_3, k) = \Omega(k^2/(\log k)^2)$ as $k \rightarrow \infty$, see [7]. In 1975 and 1977 Spencer [28, 29] reproved these results via the *Lovász Local Lemma*, and also extended them to lower bounds on $R(H, k)$ for $H \in \{K_s, C_\ell\}$, where C_ℓ denotes a cycle of length ℓ . In 1994 Krivelevich [19] further extended this to general graphs H via a new (*large-deviation based*) *alteration approach*, obtaining the lower bound

$$R(H, k) = \Omega\left((k/\log k)^{m_2(H)}\right) \quad \text{with} \quad m_2(H) := \max_{F \subseteq H} \left(\mathbb{1}_{\{v_F \geq 3\}} \frac{e_F - 1}{v_F - 2} + \mathbb{1}_{\{F=K_2\}} \frac{1}{2} \right) \quad (1)$$

as $k \rightarrow \infty$, where the implicit constants may depend on H (writing $v_F := |V(F)|$ and $e_F := |E(F)|$ for the number of vertices and edges of a graph F , and denoting by $\mathbb{1}_{\{\mathcal{E}\}}$ the indicator variable for the event \mathcal{E} , as usual). By analyzing (*semi-random*) H -free processes, in 1995 Kim [18] and in 2010 Bohman–Keevash [3] have further improved the logarithmic factors in (1) for some forbidden graphs H such as triangles K_3 , cliques K_s , and cycles C_ℓ . In 2019 Mubayi and Verstraete [24] used *pseudorandom graphs* to obtain further logarithmic improvements of (1) for cycles C_ℓ of odd length $\ell \geq 5$. However, despite considerable effort, for $H \neq K_3$ the best known lower and upper bounds on the graph Ramsey number $R(H, k)$ are still polynomial factors apart, see [3, 4, 10]. Unsurprisingly, in quest for progress research has thus stretched in several directions, including (i) variants of Ramsey numbers and (ii) further advancements of existing probabilistic proof methods.

In this paper we present a *refinement of the alteration method* for constructing H -free graphs: for suitable edge-probabilities $p = p(n)$, we show that removing all edges in H -copies of the binomial random graph $G_{n,p}$ does not significantly change the independence number, i.e., the size of the largest independent set (see Section 1.1). The main innovation is that we can allow for deleting all edges in H -copies, and not just one edge from each H -copy as in earlier alteration approaches of Erdős [7] and Krivelevich [19]. Our refinement is a natural random graph statement in its own right, and we demonstrate its usefulness via two applications to

*Faculty of Mathematics, Technion, Haifa 32000, Israel. E-mail: hguo@campus.technion.ac.il.

†Department of Mathematics, University of California, San Diego, La Jolla CA 92093, USA. E-mail: lwarnke@ucsd.edu. Supported by NSF grant DMS-1703516, NSF CAREER grant DMS-1945481, and a Sloan Research Fellowship.

online graph Ramsey games, where it enables easier analysis. In particular, for Ramsey, Paper, Scissors games and online Ramsey numbers this allows us to extend bounds of Conlon, Fox, Grinshpun and He [6] and Fox, He and Wigderson [11] to a large class of forbidden graphs H (see Section 2).

1.1 Refined alteration method

To motivate our refined alteration approach, we first review the classical alteration arguments for the Ramsey bound (1) due to Erdős [7] and Krivelevich [19]. They both use a binomial random graph $G_{n,p}$ with $n = \Theta((k/\log k)^{m_2(H)})$ vertices and edge-probability $p = \Theta((\log k)/k)$ to construct an n -vertex subgraph $G \subseteq G_{n,p}$ with two properties: (i) G is H -free, and (ii) G contains no independent set K of size k ; both properties together imply $R(H, k) > n$. Chernoff bounds suggest that the number $X_K := |E(G_{n,p}[K])|$ of edges of $G_{n,p}$ inside K is around $\binom{k}{2}p$, so for property (ii) it suffices to show that the alteration from $G_{n,p}$ to G does not remove ‘too many’ edges from each k -vertex subset K . To illustrate that this is a non-trivial task, let us consider the natural upper bound $e_H \cdot |\mathcal{H}_K|$ on the number of removed edges from K , where \mathcal{H}_K denotes the collection of all H -copies that have at least one edge inside K . For any $\delta > 0$ it turns out that $\mathbb{P}(|\mathcal{H}_K| \geq \delta \binom{k}{2}p) \geq e^{-o(k)}$ due to ‘infamous’ upper tail [16, 26] behavior (see Appendix A), which suggests that one has to rather carefully handle edges that are contained in multiple H -copies.

For triangles $H = K_3$, Erdős [7] overcame these difficulties in 1961 by a clever ad-hoc greedy alteration argument, showing that whp¹ (with high probability) the following works: If one sequentially traverses the edges of $G_{n,p}$ in any order, only accepting edges that do not create a triangle together with previously accepted edges, then the resulting ‘accepted’ subgraph $G \subseteq G_{n,p}$ satisfies the independent set property (ii), and trivially the H -free property (i). The fact that any edge-order works was exploited by Conlon, Fox, Grinshpun and He [6] and Fox, He and Wigderson [11] in their analysis of triangle-based online Ramsey games.

To handle general graphs H , Krivelevich [19] developed in 1994 an elegant alteration argument, showing that whp the following works: If one constructs $G \subseteq G_{n,p}$ by deleting all edges that are in some maximal (under inclusion) collection \mathcal{C} of edge-disjoint H -copies in $G_{n,p}$, then this (a) removes less than $X_K \approx \binom{k}{2}p$ edges from each k -vertex subset K , and (b) yields an H -free graph by maximality of \mathcal{C} , establishing both properties (ii) and (i). Unfortunately, this slick argument is hard to adapt to online Ramsey games, where one cannot foresee whether in future turns a given edge will be contained in an H -copy or not (cf. Section 2).

Our refined alteration approach overcomes the discussed difficulties, by showing that whp the desired H -free and independent set properties (i) and (ii) remain valid even if one deletes *all* edges from $G_{n,p}$ that are in some H -copy. Theorem 1 is stated for strictly 2-balanced graphs H , i.e., which satisfy $m_2(H) > m_2(F)$ for all subgraphs $F \subsetneq H$ (this class includes many graphs of interest, including cliques K_s , cycles C_ℓ , complete multipartite graphs K_{t_1, \dots, t_r} , and hypercubes Q_d). For many applications such as (1) the restriction to strictly 2-balanced graphs is immaterial, since one can often obtain the general case by forbidding a strictly 2-balanced subgraph $H_0 \subseteq H$ with $m_2(H_0) = m_2(H)$; cf. Section 2.2. Below $X_K = |E(G_{n,p}[K])|$ denotes the number of edges of $G_{n,p}$ inside K , and Y_K denotes the number of edges in $E(G_{n,p}[K])$ that are in some H -copy of $G_{n,p}$.

Theorem 1 (Main alteration tool). *Let H be a strictly 2-balanced graph. Then, for any fixed $\delta > 0$, the following holds for all sufficiently large $C \geq C_0 = C_0(\delta, H)$ and sufficiently small $0 < c \leq c_0 = c_0(C, \delta, H)$. The random graph $G_{n,p}$ with $n := \lfloor c(k/\log k)^{m_2(H)} \rfloor$ vertices and edge-probability $p := C(\log k)/k$ whp satisfies $Y_K \leq \delta \binom{k}{2}p$ and $X_K \geq (1 - \delta) \binom{k}{2}p$ for all k -vertex subsets K of $G_{n,p}$.*

As discussed, our refined alteration method constructs $G \subseteq G_{n,p}$ by deleting all edges that are in some H -copy of $G_{n,p}$, so that G trivially satisfies the H -free property (i). For fixed $\delta \in (0, 1/2)$ and suitable n, p , our main tool Theorem 1 then ensures that the following holds whp: for any k -vertex subset K of G we have

$$|E(G[K])| = X_K - Y_K \geq (1 - 2\delta) \binom{k}{2}p > 0,$$

which implies that G also satisfies the independent set property (ii). To put this into context, we remark that the size of the largest independent set of $G_{n,p}$ whp satisfies $\alpha(G_{n,p}) = \Theta(\log(np)/p) = \Theta(k)$ when $m_2(H) > 1$, see [15, Section 7.1], which implies that whp $\Theta(k) = \alpha(G_{n,p}) \leq \alpha(G) < k$. The conceptual crux is thus that deleting all edges in H -copies does not significantly change the independence number of $G_{n,p}$ (earlier alteration approaches [7, 19] only gave such guarantees after deleting a carefully chosen subset of these edges).

¹In this paper *whp* (with high probability) always means with probability tending to one as $k \rightarrow \infty$.

As we shall see in Section 2, variants of the above-discussed alteration argument carry over to certain online Ramsey games (where arbitrary deletion of edges in H -copies will be key). We believe that the upper bound on Y_K from Theorem 1 will also be useful in other contexts, since the classical alteration method has become a standard tool in many applications, including coloring problems for graphs with forbidden substructures [21, 12, 2], Ramsey number variants [19, 30, 6, 11], induced bipartite triangle-free graphs [8, 23], algorithmic approximation ratios [20], and semi-random constructions [18, 13, 14]. Further extensions of our refined alteration method are discussed in Section 4.

2 Applications to online Ramsey games

In this section we demonstrate the usefulness of our refined alteration method via two applications to online graph Ramsey games. Here it will be crucial that we can allow for arbitrary deletion of edges in H -copies, which enables easier analysis in the online setting (by similar reasoning as in the classical ‘offline’ setting).

2.1 Ramsey, Paper, Scissors Game

Our first application concerns the *Ramsey, Paper, Scissors game* that was introduced by Fox, He and Wigderson [11]. For a graph H , this is a game between two players, Proposer and Decider, that starts with a finite set $V = \{1, 2, \dots, n\}$ of n isolated vertices. In each turn, Proposer proposes a pair of non-adjacent vertices from V , and Decider simultaneously decides whether or not to add it as an edge to the current graph (without knowing which pair is proposed). Proposer cannot propose vertex-pairs that would form a copy of H together the current graph, nor vertex-pairs that have been proposed before. The *RPS number* $\text{RPS}(H, n)$ is defined² as the largest number k for which Proposer can guarantee that, with probability at least $1/2$ (regardless of Decider’s strategy), the final graph has an independent set of size k . Our refined alteration method enables us to prove the following upper bound on $\text{RPS}(H, n)$ for all strictly 2-balanced graphs H .

Theorem 2 (Ramsey, Paper, Scissors Game). *If H is a strictly 2-balanced graph, then the RPS number satisfies $\text{RPS}(H, n) = O(n^{1/m_2(H)} \log n)$ as $n \rightarrow \infty$, where the implicit constant may depend on H .*

For all strictly 2-balanced graphs H , Theorem 2 gives the best known upper bounds for RPS numbers. For s -vertex cliques we obtain $\text{RPS}(K_s, n) = O(n^{2/(s+1)} \log n)$, which generalizes the upper bound part of the $\text{RPS}(K_3, n) = \Theta(\sqrt{n} \log n)$ result of Fox, He and Wigderson [11] for triangles.

The following proof of Theorem 2 demonstrates that for Ramsey, Paper, Scissors games, our refined alteration approach leads to conceptually simple upper bound proofs.

Proof of Theorem 2. For $\delta := 1/4$ we choose $C > 0$ large enough and then $c > 0$ small enough so that Theorem 1 applies to $G_{n,p}$ with $n := \lfloor c(k/\log k)^{m_2(H)} \rfloor$ and $p := C(\log k)/k$. We shall analyze the following Decider strategy: in each turn Decider accepts the (unknown) proposed vertex-pair as an edge independently with probability p . To prove $\text{RPS}(H, n) < k = O(n^{1/m_2(H)} \log n)$ as $n \rightarrow \infty$ (where the implicit constant depends on H), it suffices to show that the resulting final graph has whp no independent set of size k .

Turning to details, let G denote the resulting final graph at the end of the game, i.e., which contains all accepted edges. Since all edges that do not create H -copies are eventually proposed, there is a natural coupling between $G_{n,p}$ and G which satisfies the following two properties: (a) that $E(G) \subseteq E(G_{n,p})$, and (b) that every edge in $E(G_{n,p}) \setminus E(G)$ is contained in an H -copy of $G_{n,p}$. Invoking Theorem 1, it follows that this coupling satisfies the following whp: for any k -vertex subset K of G we have

$$|E(G[K])| \geq X_K - Y_K \geq (1 - 2\delta) \binom{k}{2} p = \frac{1}{2} \binom{k}{2} p > 0,$$

which implies that the final graph G has whp no independent set of size k , so that the desired upper bound $\text{RPS}(H, n) < k = O(n^{1/m_2(H)} \log n)$ follows (as discussed above). \square

²For imperfect-information games such as Ramsey, Paper, Scissors (both players make simultaneous moves) one usually considers randomized strategies, see [25, pp. 14, 169], motivating why the definition of $\text{RPS}(H, n)$ includes probability of winning.

2.2 Online Ramsey Game

Our second application concerns the widely-studied *online Ramsey game* (see, e.g., [1, 22, 17, 5, 6]) that was introduced independently by Beck [1] and Kurek–Ruciński [22]. This is a game between two players, Builder and Painter, that starts with an infinite set $V = \{1, 2, \dots\}$ of isolated vertices. In each turn, Builder places an edge between two non-adjacent vertices from V , and Painter immediately colors it either red or blue. The *online Ramsey number* $\tilde{r}(H, k)$ is defined as the smallest number of turns N that Builder needs to guarantee the existence of either a red copy of H or a blue copy of K_k (regardless of Painter’s strategy). Our refined alteration method enables us to prove the following lower bound on $\tilde{r}(H, k)$, which, up to logarithmic factors, is about k times the best-known general lower bound for the usual Ramsey number $R(H, k)$; cf. (1).

Theorem 3 (Online Ramsey Game). *If H is a graph with $e_H \geq 1$ edges, then the online Ramsey number satisfies $\tilde{r}(H, k) = \Omega(k \cdot (k/\log k)^{m_2(H)})$ as $k \rightarrow \infty$, where the implicit constant may depend on H .*

For general graphs H , Theorem 3 gives the best known lower bounds for online Ramsey numbers. For s -vertex cliques we obtain $\tilde{r}(K_s, k) = \Omega(k^{(s+3)/2}/(\log k)^{(s+1)/2})$, which generalizes the lower bound of Conlon, Fox, Grinshpun and He [6, Theorem 4] for triangles, and improves [6, Corollary 3] for small clique sizes s .

The following proof of Theorem 3 demonstrates that for Online Ramsey games, our refined alteration approach again leads to proofs that conceptually mimic the usual reasoning from the classical offline setting. To this end we shall analyze a semi-random Painter strategy, whose main ideas we now outline. The default color of an edge is blue. But if an edge is placed between vertices of sufficiently ‘high’ degree, it does the following independently with probability p : it colors the edge red, unless this would create a red H -copy, in which case the edge is still colored blue (i.e., the red-coloring attempt is rejected). By construction there are no red H -copies, and it turns out that blue cliques K_k can only appear inside the growing set U of high-degree vertices (since vertices in K_k must have ‘high’ degree). Ignoring a number of technicalities, after the first $N = \Theta(k \cdot (k/\log k)^{m_2(H)})$ turns, we are able to establish that, inside each k -vertex subset $K \subseteq U$, the following holds: (i) the number of ‘rejected’ red-coloring attempts is at most Y_K , and (ii) the total number of red-coloring attempts is at least $X_K^* \sim \text{Bin}(\frac{1}{2}\binom{k}{2}, p)$. Using our refined alteration approach and Chernoff Bounds we can then show that, whp, every k -vertex subset $K \subseteq U$ contains at least $X_K^* - Y_K > 0$ red edges, which prevents blue cliques K_k and thus gives $\tilde{r}(H, k) > N = \Theta(k \cdot (k/\log k)^{m_2(H)})$. The technical details of the following proof are complicated by the fact that the online Ramsey game is played on an infinite set $V = \{1, 2, \dots\}$ of vertices, which requires special care in the coupling and union bound arguments.

Proof of Theorem 3. For convenience we first suppose that H is strictly 2-balanced. For $\delta := 1/8$ we choose $C \geq 64e_H$ large enough and then $c > 0$ small enough so that Theorem 1 applies to $G_{n,p}$ with $n := \lfloor c(k/\log k)^{m_2(H)} \rfloor$ and $p := C(\log k)/k$. Set $L := \lfloor (k-1)/4 \rfloor$. At any moment of the game, we define $U \subseteq V$ as the set of all vertices that, in the current graph, are adjacent to at least L edges placed by builder (to clarify: the growing vertex set U is updated at the end of each turn).

We shall analyze the following Painter strategy: Painter’s default color is blue, but if an edge $e = \{x, y\}$ is placed inside U , then Painter does the following independently with probability p (\star): it colors the edge e red, unless this would create a red H -copy (\dagger), in which case the edge e is still colored blue. By construction there are no red H -copies, and blue cliques K_k can only appear inside U (since all vertices in copy of K_k must be adjacent to at least $k-1 > L$ vertices). To prove $\tilde{r}(H, k) > N := \lfloor L \cdot n/2 \rfloor = \Omega(k \cdot (k/\log k)^{m_2(H)})$ as $k \rightarrow \infty$ (with implicit constants depending on H), by the usual reasoning it remains to show that after N turns there are whp no blue cliques K_k inside U . Let \mathcal{K} denote the collection of all k -vertex subsets $K \subseteq U$ after N turns. The plan is to show that, inside each vertex set $K \in \mathcal{K}$ that can become a blue clique K_k , there are more red-coloring attempts (\star) than ‘rejected’ red-coloring attempts (\dagger), which enforces a red edge inside K .

Turning to details, note that $|U| \leq 2N/L \leq n$ during the first N turns. Using the order in which vertices enter U (breaking ties using lexicographic order), at any moment during the first N turns we thus obtain an injection $\Phi : U \mapsto \{1, \dots, n\} = V(G_{n,p})$. After N turns, we abbreviate this injection by Φ_N , and write $\Phi_N(K) := \{\Phi_N(v) : v \in K\}$. Define \mathcal{B}_K as the event that, during the first N turns, the number of ‘rejected’ red-coloring attempts (\dagger) inside K is at most $\frac{1}{8}\binom{k}{2}p$. There is a natural turn-by-turn inductive coupling between $G_{n,p}$ and Painter’s strategy, where the red-coloring attempt (\star) occurs if $\Phi(e) := \{\Phi(x), \Phi(y)\}$ is an edge of $G_{n,p}$. A moments thought reveals that, during the first N turns, under this coupling the total number of ‘rejected’ red-colorings (\dagger) inside $K \in \mathcal{K}$ is at most $Y_{\Phi_N(K)}$ defined with respect to $G_{n,p}$ (since (\dagger) can only happen when a red-coloring of $e \subseteq K$ creates a red H -copy, which under the coupling implies

that $\Phi(e) \subseteq \Phi(K)$ is contained in an H -copy of $G_{n,p}$). Applying Theorem 1 with $\delta = 1/8$ to $G_{n,p}$, using the described coupling and $|\Phi_N(K)| = |K| = k$ it then follows that, whp, the event \mathcal{B}_K occurs for all $K \in \mathcal{K}$.

Intuitively, we shall next show that, for all k -vertex sets $K \in \mathcal{K}$ that contain $\binom{k}{2}$ edges (a prerequisite for having a blue clique K_k inside K), the total number of red-coloring attempts (\star) inside K is at least $\frac{1}{4}\binom{k}{2}p$. To make this precise, define \mathcal{T}_K as the event that builder places less than $\binom{k}{2}$ edges inside K during the first N turns. Let X_K^* denote the number of red-coloring attempts (\star) inside K during the first N turns, and define \mathcal{A}_K as the event that $X_K^* \geq \frac{1}{4}\binom{k}{2}p$. Let \mathcal{K}' denote the collection of all k -vertex subsets $K' \subseteq V(G_{n,p})$. Since Φ_N defines an injection from \mathcal{K} to \mathcal{K}' , writing $\Phi_N^{-1}(K') := \{v \in V : \Phi_N(v) \in K'\}$ it follows that

$$\mathbb{P}(\neg\mathcal{A}_K \cap \neg\mathcal{T}_K \text{ for some } K \in \mathcal{K}) \leq \sum_{K' \in \mathcal{K}'} \mathbb{P}\left(X_{\Phi_N^{-1}(K')}^* \leq \frac{1}{4}\binom{k}{2}p \text{ and } \neg\mathcal{T}_{\Phi_N^{-1}(K')}\right). \quad (2)$$

Fix $K' \in \mathcal{K}'$, and set $K := \Phi_N^{-1}(K')$. Note that, by checking in each turn for red-coloring attempts (\star) inside $\Phi^{-1}(K') := \{v \in V : \Phi(v) \in K'\}$, we can determine X_K^* without knowing Φ_N^{-1} in advance. Furthermore, since every vertex is adjacent to at most L vertices before entering U , the event $\neg\mathcal{T}_K$ implies that during the first N turns at least $\binom{k}{2} - |K| \cdot L \geq \frac{1}{2}\binom{k}{2}$ red-coloring attempts (\star) happen inside K , each of which is (conditional on the history) successful with probability p . It follows that X_K^* stochastically dominates a binomial random variable $Z \sim \text{Bin}(\lceil \frac{1}{2}\binom{k}{2} \rceil, p)$, unless the event \mathcal{T}_K occurs. Noting $kp = C \log k \geq 64e_H \log k$ and $n \ll k^{e_H}$, by invoking standard Chernoff bounds (such as [15, Theorem 2.1]) it then follows that

$$\mathbb{P}\left(X_{\Phi_N^{-1}(K')}^* \leq \frac{1}{4}\binom{k}{2}p \text{ and } \neg\mathcal{T}_{\Phi_N^{-1}(K')}\right) \leq \mathbb{P}\left(Z \leq \frac{1}{4}\binom{k}{2}p\right) \leq \exp\left(-\binom{k}{2}p/16\right) \ll k^{-e_H k} \ll n^{-k}. \quad (3)$$

Combining (2)–(3) with $|\mathcal{K}'| \leq n^k$, we readily infer that, whp, the event $\mathcal{A}_K \cup \mathcal{T}_K$ occurs for all $K \in \mathcal{K}$.

To sum up, the following holds whp after N turns: every k -vertex subset $K \subseteq U$ contains either (a) at least $\frac{1}{4}\binom{k}{2}p - \frac{1}{8}\binom{k}{2}p = \frac{1}{4}\binom{k}{2}p > 0$ red edges, or (b) less than $\binom{k}{2}$ edges in total. Both possibilities prevent a blue clique K_k inside K , and so the desired lower bound $\tilde{r}(H, k) > N$ follows (as discussed above).

Finally, in the remaining case where H is not strictly 2-balanced, we pick a minimal subgraph $H_0 \subsetneq H$ with $m_2(H_0) = m_2(H)$. It is straightforward to check that, by construction, H_0 is strictly 2-balanced. Furthermore, since any H_0 -free graph is also H -free, we also have $\tilde{r}(H, k) \geq \tilde{r}(H_0, k)$. Repeating the above proof with H replaced by H_0 then gives the claimed lower bound on $\tilde{r}(H, k)$. \square

It would be interesting to investigate whether Theorem 3 can be improved if one replaces our $G_{n,p}$ based alteration approach by an H -free process [3] based approach or semi-random variants thereof [18, 13].

3 Proof of main alteration tool Theorem 1

In this section we prove the main tool Theorem 1 of our refined alteration method. The main difficulty is the desired upper bound on Y_K , which denotes the number of edges in $E(G_{n,p}[K])$ that are in some H -copy of $G_{n,p}$. Here our core proof strategy is to approximate Y_K by more tractable auxiliary random variables, inspired by ideas from [16, 33, 32, 27]. In particular, we expect that the main contribution to Y_K should come from H -copies that share exactly two vertices and one edge with K ; in the below proof we will denote the collection of such ‘good’ H -copies by \mathcal{H}_K^* . Note that when multiple good H -copies from \mathcal{H}_K^* contain some common edge f inside K , they together only contribute one edge to Y_K . It follows that, by arbitrarily selecting one ‘representative’ copy $H_f \in \mathcal{H}_K^*$ for each relevant edge f , we should obtain a sub-collection $\mathcal{H} \subseteq \mathcal{H}_K^*$ of good H -copies with $|\mathcal{H}| \approx Y_K$. The H -copies in \mathcal{H} share no edges inside K by construction, and it turns out that all other types of edge-overlaps are ‘rare’, i.e., make a negligible contribution to Y_K . We thus expect that there is an edge-disjoint sub-collection $\mathcal{H}' \subseteq \mathcal{H} \subseteq \mathcal{H}_K^*$ of good H -copies with $|\mathcal{H}'| \approx |\mathcal{H}| \approx Y_K$, and here the crux is that the upper tail of $|\mathcal{H}'|$ is much easier to estimate than the upper tail of Y_K (see Claim 5 below). The following proof implements a rigorous variant of the above-discussed heuristic ideas for bounding Y_K .

Proof of the Y_K bound of Theorem 1. Noting that the claimed bounds are trivial when $m_2(H) \leq 1$ (since then there are no k -vertex subsets K in $G_{n,p}$ due to $n \ll k$), we may henceforth assume $m_2(H) > 1$.

Fix a k -vertex set K . Let \mathcal{H}_K denote the collection of all H -copies in $G_{n,p}$ that have at least one edge inside K , and let $\mathcal{H}_K^* \subseteq \mathcal{H}_K$ denote the sub-collection of H -copies that moreover share exactly two vertices with K . Let \mathcal{I}_K denote a size-maximal collection of edge-disjoint $H \in \mathcal{H}_K^*$. Clearly $|\mathcal{I}_K| \leq Y_K$,

and Claim 4 below establishes a related upper bound. Let \mathcal{T}_K denote a size-maximal collection of edge-disjoint $H \in \mathcal{H}_K \setminus \mathcal{H}_K^*$. Let \mathcal{P}_K denote a size-maximal collection of edge-disjoint $H_1 \cup H_2$ with distinct $H_1, H_2 \in \mathcal{H}_K^*$ that satisfy $|E(H_1) \cap E(H_2)| \geq 1$ and $V(H_1) \cap K \neq V(H_2) \cap K$. Let $\Delta_{H,f}$ denote the number of H -copies in $G_{n,p}$ that contain the edge f , and define Δ_H as the maximum of $\Delta_{H,f}$ over all $f \in E(K_n)$.

Claim 4. *We have $Y_K \leq |\mathcal{I}_K| + 2e_H^2(|\mathcal{T}_K| + |\mathcal{P}_K|)\Delta_H$.*

Proof of Claim 4. We divide the H -copies in \mathcal{H}_K into two disjoint groups: those which share at least one edge with some $H \in \mathcal{T}_K$ or $H_1 \cup H_2 \in \mathcal{P}_K$, and those which do not; we denote these two groups by \mathcal{H}_1 and \mathcal{H}_2 , respectively. For $j \in \{1, 2\}$, let \mathcal{E}_j denote the collection of edges from K that are contained in at least one H -copy from \mathcal{H}_j . Note that $Y_K \leq |\mathcal{E}_1| + |\mathcal{E}_2|$ and $|\mathcal{E}_1| \leq e_H |\mathcal{H}_1| \leq e_H \cdot (e_H |\mathcal{T}_K| + 2e_H |\mathcal{P}_K|)\Delta_H$.

Gearing up towards bounding $|\mathcal{E}_2|$, recall that \mathcal{H}_2 contains all H -copies in \mathcal{H}_K that are edge-disjoint from all $H \in \mathcal{T}_K$ and $H_1 \cup H_2 \in \mathcal{P}_K$. By size-maximality of \mathcal{T}_K , it follows that $\mathcal{H}_2 \subseteq \mathcal{H}_K^*$ (since any H -copy in \mathcal{H}_2 that shares more than two vertices with K could be added to \mathcal{T}_K , contradicting maximality). By size-maximality of \mathcal{P}_K , we then infer the following property (\bullet) of \mathcal{H}_2 : any two distinct H -copies in \mathcal{H}_2 are edge-disjoint, unless they both intersect K in the same two vertices (since otherwise we could add them to \mathcal{P}_K , contradicting maximality). For each $f \in \mathcal{E}_2 \subseteq \binom{K}{2}$ we now arbitrarily select one H -copy from \mathcal{H}_2 that contains f . By property (\bullet) and the size-maximality of \mathcal{I}_K , this yields a sub-collection $\mathcal{H}'_2 \subseteq \mathcal{H}_2 \subseteq \mathcal{H}_K^*$ of edge-disjoint H -copies satisfying $|\mathcal{E}_2| = |\mathcal{H}'_2| \leq |\mathcal{I}_K|$, and the claim follows. \square

The remaining upper tail bounds for $|\mathcal{I}_K|$, $|\mathcal{T}_K|$, $|\mathcal{P}_K|$ and Δ_H hinge on the following four key estimates. First, $m_2(H) > 1$ and strictly 2-balancedness of H imply $m_2(H) = (e_H - 1)/(v_H - 2)$, so that

$$n^{v_H-2} p^{e_H-1} = (np^{m_2(H)})^{v_H-2} \leq (cC^{m_2(H)})^{v_H-2}. \quad (4)$$

Second, $n = k^{m_2(H)-o(1)}$ and $m_2(H) > 1$ imply that there is $\tau = \tau(H) > 0$ such that

$$\frac{k}{n} = k^{1-m_2(H)+o(1)} \ll k^{-\tau} / \log k. \quad (5)$$

Third, using $p = k^{-1+o(1)}$ and strictly 2-balancedness of H (implying that $(e_J - 1)/(v_J - 2) < m_2(H)$ for all $J \subsetneq H$ with $e_J \geq 2$), it follows that there is $\gamma = \gamma(H) > 0$ such that

$$n^{v_J-2} p^{e_J-1} = (np^{(e_J-1)/(v_J-2)})^{v_J-2} \gg k^\gamma \quad \text{for all } J \subsetneq H \text{ with } e_J \geq 2. \quad (6)$$

The below-claimed fourth estimate can be traced back to Erdős and Tetali [9]; we include an elementary proof for self-containedness (see [33, Section 2] for related estimates that also allow for overlapping edge-sets).

Claim 5. *Let \mathcal{S} be a collection of edge-subsets from $E(K_n)$. Set $\mu := \sum_{\beta \in \mathcal{S}} \mathbb{P}(\beta \subseteq E(G_{n,p}))$. Define Z as the largest number of disjoint edge-sets from \mathcal{S} that are present in $G_{n,p}$. Then $\mathbb{P}(Z \geq x) \leq (e\mu/x)^x$ for all $x > \mu$.*

Proof of Claim 5. Set $s := \lceil x \rceil \geq 1$. Exploiting edge-disjointness and $s! \geq (s/e)^s$, it follows that

$$\mathbb{P}(Z \geq x) \leq \sum_{\substack{\{\beta_1, \dots, \beta_s\} \subseteq \mathcal{S}: \\ \text{all edge-disjoint}}} \underbrace{\mathbb{P}(\beta_1 \cup \dots \cup \beta_s \subseteq E(G_{n,p}))}_{=\prod_{1 \leq i \leq s} \mathbb{P}(\beta_i \subseteq E(G_{n,p}))} \leq \frac{1}{s!} \left(\sum_{\beta \in \mathcal{S}} \mathbb{P}(\beta \subseteq E(G_{n,p})) \right)^s \leq (e\mu/s)^s,$$

which completes the proof by noting that the function $s \mapsto (e\mu/s)^s$ is decreasing for positive $s \geq \mu$. \square

We are now ready to bound the probability that $|\mathcal{I}_K|$ is large. Since H is strictly 2-balanced, it contains no isolated vertices and thus is uniquely determined by its edge-set. This enables us to apply Claim 5 to $|\mathcal{I}_K| = Z$ (as \mathcal{I}_K is a size-maximal collection of edge-disjoint H -copies from \mathcal{H}_K^*). Using estimate (4), it is routine to see that, for small enough $c \leq c_0(C, \delta, H)$, the associated parameter μ from Claim 5 satisfies

$$\mu \leq O(k^2 n^{v_H-2} \cdot p^{e_H}) \leq \binom{k}{2} p \cdot \Theta(n^{v_H-2} p^{e_H-1}) \leq \frac{\delta}{2e^2} \binom{k}{2} p. \quad (7)$$

Noting $\delta k p = \delta C \log k$ and $n \ll k^{e_H}$, now Claim 5 (with $Z = |\mathcal{I}_K|$) implies that, for large enough $C \geq C_0(\delta, H)$, we have

$$\mathbb{P}\left(|\mathcal{I}_K| \geq \frac{\delta}{2} \binom{k}{2} p\right) \leq \left(\frac{e\mu}{\frac{\delta}{2} \binom{k}{2} p}\right)^{\frac{\delta}{2} \binom{k}{2} p} \leq e^{-\frac{\delta}{2} \binom{k}{2} p} \ll k^{-e_H k} \ll n^{-k}. \quad (8)$$

Next, we similarly use Claim 5 to bound the probability that $|\mathcal{T}_K|$ is large. For the associated parameter μ we shall proceed similar to (7): using estimates (4)–(5), for small enough $c \leq c_0(C, \delta, H)$ we obtain

$$\mu \leq O(k^3 n^{v_H-3} \cdot p^{e_H}) \leq \binom{k}{2} p \cdot \frac{k}{n} \cdot \Theta(n^{v_H-2} p^{e_H-1}) \leq k^{-\tau} \cdot \frac{\delta}{e} \binom{k}{2} p / \log k. \quad (9)$$

With similar considerations as for (8), for large enough $C \geq C_0(\tau, \delta, H)$ Claim 5 (with $Z = |\mathcal{T}_K|$) then yields

$$\mathbb{P}\left(|\mathcal{T}_K| \geq \delta \binom{k}{2} p / \log k\right) \leq k^{-\tau \delta \binom{k}{2} p / \log k} = e^{-\tau \delta \binom{k}{2} p} \ll k^{-e_H k} \ll n^{-k}. \quad (10)$$

We shall analogously use Claim 5 to bound the probability that $|\mathcal{P}_K|$ is large. For the associated parameter μ , the basic idea is to distinguish all possible subgraphs $J \subsetneq H$ in which the relevant $H_1, H_2 \in \mathcal{H}_K^*$ can intersect. Also taking into account the number of vertices which H_1 and H_2 have inside K , i.e., $|V(H_1) \cup V(H_2) \cap K| \in \{3, 4\}$, by definition of \mathcal{P}_K it now follows via estimates (4)–(6) that

$$\begin{aligned} \mu &\leq \sum_{J \subsetneq H: e_J \geq 1} O\left(k^3 n^{2(v_H-2)-(v_J-1)} \cdot p^{2e_H-e_J} + k^4 n^{2(v_H-2)-v_J} \cdot p^{2e_H-e_J}\right) \\ &\leq \binom{k}{2} p \cdot \left[\frac{k}{n} + \left(\frac{k}{n}\right)^2\right] \cdot \sum_{J \subsetneq H: e_J \geq 1} \frac{\Theta((n^{v_H-2} p^{e_H-1})^2)}{n^{v_J-2} p^{e_J-1}} \leq k^{-\tau} \cdot \frac{\delta}{e} \binom{k}{2} p / \log k. \end{aligned} \quad (11)$$

(To clarify: in (11) above we used that (6) implies $n^{v_J-2} p^{e_J-1} \geq 1$ for all $J \subsetneq H$ with $e_J \geq 1$.) Similarly to inequalities (8) and (10), for large enough $C \geq C_0(\tau, \delta, H)$ now Claim 5 (with $Z = |\mathcal{P}_K|$) yields

$$\mathbb{P}\left(|\mathcal{P}_K| \geq \delta \binom{k}{2} p / \log k\right) \leq k^{-\tau \delta \binom{k}{2} p / \log k} = e^{-\tau \delta \binom{k}{2} p} \ll k^{-e_H k} \ll n^{-k}. \quad (12)$$

Finally, combining (8), (10) and (12) with Claim 4, a standard union bound argument gives

$$\mathbb{P}\left(Y_K \geq \delta \binom{k}{2} p \cdot \left(\frac{1}{2} + 4e_H^2 \Delta_H / \log k\right) \text{ for some } k\text{-vertex set } K\right) \leq \binom{n}{k} \cdot o(n^{-k}) = o(1). \quad (13)$$

To complete the proof of the Y_K bound, it thus remains to show that, for small enough $c \leq c_0(C, H)$, we have

$$\mathbb{P}\left(\Delta_H \geq (\log k) / (8e_H^2)\right) = o(1). \quad (14)$$

Using (4), (6) and $n \ll k^{e_H}$, this upper tail estimate for $\Delta_H = \max_f \Delta_{H,f}$ follows routinely from standard concentration inequalities such as [32, Theorem 32], but we include an elementary proof for self-containedness (based on ideas from [31, 33]). Turning to the proof of (14), let $\Delta_{H,f,g}$ denote the number of H -copies in $G_{n,p}$ that contain the edges $\{f, g\}$, and define $\Delta_H^{(2)}$ as the maximum of $\Delta_{H,f,g}$ over all distinct $f, g \in E(K_n)$. We call an r -tuple (H_1, \dots, H_r) of H -copies an (r, f, g) -star if each H_j contains the edges $\{f, g\}$ and satisfies $H_j \not\subseteq H_1 \cup \dots \cup H_{j-1}$. Define $Z_{r,f,g}$ as the number of (r, f, g) -stars (H_1, \dots, H_r) that are present in $G_{n,p}$. Summing over all $(r+1, f, g)$ -stars (H_1, \dots, H_{r+1}) , by noting that the intersection of H_{r+1} with $F_r := H_1 \cup \dots \cup H_r$ is isomorphic to some proper subgraph $J \subsetneq H$ containing at least $e_J \geq 2$ edges, using estimates (4) and (6) it then is routine to see that, for $1 \leq r \leq r_0 := 1 + \lceil (v_H e_H + 4e_H) / \gamma \rceil$, we have

$$\begin{aligned} \mathbb{E} Z_{r+1,f,g} &= \sum_{(H_1, \dots, H_{r+1})} p^{e_{H_1 \cup \dots \cup H_{r+1}}} = \sum_{(H_1, \dots, H_r)} p^{e_{F_r}} \sum_{H_{r+1}} p^{e_H - e_{H_{r+1} \cap F_r}} \\ &\leq \sum_{(H_1, \dots, H_r)} p^{e_{F_r}} \cdot \sum_{J \subsetneq H: e_J \geq 2} O\left((v_H r)^{v_J} n^{v_H - v_J} \cdot p^{e_H - e_J}\right) \leq \mathbb{E} Z_{r,f,g} \cdot k^{-\gamma}. \end{aligned}$$

Since trivially $\mathbb{E} Z_{1,f,g} = O(n^{v_H})$, using $n \ll k^{e_H}$ we infer $\mathbb{E} Z_{r_0,f,g} \leq k^{v_H e_H - (r_0-1)\gamma} \leq k^{-4e_H} \ll n^{-4}$. Consider a maximal length (r, f, g) -star (H_1, \dots, H_r) in $G_{n,p}$, and note that in $G_{n,p}$ any H -copy containing the edges $\{f, g\}$ is completely contained in $H_1 \cup \dots \cup H_r$ (by length maximality), so that $\Delta_{H,f,g} \leq (e_H r)^{e_H}$ holds (using that H is uniquely determined by its edge-set). For $D := (e_H r_0)^{e_H}$ it follows that

$$\mathbb{P}\left(\Delta_H^{(2)} \geq D\right) \leq \sum_{f \neq g} \mathbb{P}(\Delta_{H,f,g} \geq D) \leq \sum_{f \neq g} \mathbb{P}(Z_{r_0,f,g} \geq 1) \leq \sum_{f \neq g} \mathbb{E} Z_{r_0,f,g} \leq \binom{n}{2}^2 \cdot o(n^{-4}) = o(1). \quad (15)$$

With an eye on $\Delta_{H,f}$, let \mathcal{H}_f denote the collection of all H -copies in K_n that contain the edge f . We pick a subset $\mathcal{I} \subseteq \mathcal{H}_f$ of H -copies in $G_{n,p}$ that is size-maximal subject to the restriction that all H -copies are edge-disjoint after removing the common edge f . For any $H' \in \mathcal{H}_f$, note that in $G_{n,p}$ there are a total of at most $e_H \Delta_H^{(2)}$ copies of H that share f and at least one additional edge with H' . Hence $\Delta_{H,f} \geq (\log k)/(8e_H^2)$ and $\Delta_H^{(2)} \leq D$ imply $|\mathcal{I}| \geq \lceil (\log k)/A \rceil =: z$ for $A := 8e_H^3 D$ (by maximality of \mathcal{I}). As the union of all H -copies in \mathcal{I} contains exactly $1 + (e_H - 1)|\mathcal{I}|$ edges, using $\binom{m}{z} \leq (em/z)^z$ and $|\mathcal{H}_f| = O(n^{v_H-2})$ it follows that

$$\mathbb{P}\left(\Delta_{H,f} \geq (\log k)/(8e_H^2) \text{ and } \Delta_H^{(2)} \leq D\right) \leq \binom{|\mathcal{H}_f|}{z} \cdot p^{1+(e_H-1)z} \leq \left(\frac{O(n^{v_H-2} p^{e_H-1})}{z}\right)^z. \quad (16)$$

Using estimate (4), for small enough $c \leq c_0(A, C, H)$ the right-hand side of (16) is at most $(\log k)^{-(\log k)/A} \ll k^{-2e_H}$. Recalling $n \ll k^{e_H}$, by taking a union bound over all edges $f \in E(K_n)$ it then follows that

$$\mathbb{P}\left(\Delta_H \geq (\log k)/(8e_H^2) \text{ and } \Delta_H^{(2)} \leq D\right) \leq \binom{n}{2} \cdot o(k^{-2e_H}) = o(1), \quad (17)$$

which together with (15) completes the proof of estimate (14) and thus the Y_K bound of Theorem 1. \square

The above proof of (14) can easily be sharpened to $\mathbb{P}(\Delta_H \geq B(\log k)/\log \log k) = o(1)$ for a suitable constant $B = B(H) > 0$, see (16)–(17). Together with the proof of (13) and $|\mathcal{I}_K| \leq Y_K$, this implies that whp $Y_K = |\mathcal{I}_K| + o(\delta \binom{k}{2} p)$ for all k -vertex subsets K , which suggests that Y_K is well-approximated by $|\mathcal{I}_K|$.

To complete the proof of Theorem 1, it remains to give the routine Chernoff bound based proof of the X_K bound of Theorem 1, where $X_K = |E(G_{n,p}[K])|$ denotes the number of edges of $G_{n,p}$ inside K .

Proof of the X_K bound of Theorem 1. Noting $X_K \sim \text{Bin}\left(\binom{k}{2}, p\right)$ as well as $\delta^2 kp = \delta^2 C \log k$ and $n \ll k^{e_H}$, by standard Chernoff bounds (such as [15, Theorem 2.1]) it follows, for large enough $C \geq C_0(\delta, H)$, that

$$\mathbb{P}(X_K \leq (1 - \delta) \binom{k}{2} p) \leq \mathbb{1}_{\{\delta \in (0,1]\}} \exp\left(-\delta^2 \binom{k}{2} p/2\right) \ll k^{-e_H k} \ll n^{-k}.$$

Taking a union bound over all k -vertex subsets K completes the proof of the X_K bound of Theorem 1. \square

4 Extensions

In applications of the alteration method outlined in Section 1.1, it often is beneficial to keep track of further properties of the resulting H -free n -vertex graph $G \subseteq G_{n,p}$, including vertex-degrees and the number of edges (see, e.g., [8, Section 3], [2, Section 2], and [23, Section 5.1]). Using the arguments and intermediate results from Section 3, oftentimes it is routine to show that G resembles a random graph $G_{n,p}$ in many ways. For example, with standard results for $G_{n,p}$ in mind, the following simple lemma intuitively implies that whp the resulting G is approximately np regular, with about $\binom{n}{2} p$ edges (note that $np \gg 1$ when $m_2(H) > 1$).

Lemma 6. *Let H be a strictly 2-balanced graph with $m_2(H) > 1$. Define Y as the number of H -copies in $G_{n,p}$, and define Y_v as the number of H -copies in $G_{n,p}$ that contain the vertex v . For any fixed $\delta > 0$, the following holds for all sufficiently large $C \geq C_0 = C_0(\delta, H)$ and sufficiently small $0 < c \leq c_0 = c_0(C, \delta, H)$. Setting n and p as in Theorem 1, whp $G_{n,p}$ satisfies $Y_v \leq \delta np$ for all vertices v , and $Y \leq \delta \binom{n}{2} p$.*

Proof. Since $m_2(H) > 1$ implies $v_H \geq 3$, noting $Y = \sum_{v \in [n]} Y_v/v_H$ it suffices to prove the claimed bounds on the Y_v . Fix a vertex v . Similar to estimate (7), using (4) it is standard to see that the expected number of H -copies containing v is at most $\mu \leq O(n^{v_H-1} p^{e_H}) \leq \frac{\delta}{c^2} np$ for small enough $c \leq c_0(C, \delta, H)$. Furthermore, if $\Delta_H \leq (\log k)/(8e_H^2)$ holds (see (14) in Section 3), then any H -copy edge-intersects a total of at most $e_H \cdot \Delta_H < \log k$ many H -copies, say. Applying the upper tail inequality [13, Theorem 15] instead of Claim 5, using $\delta np = \delta c C k^{m_2(H)-1-o(1)} \gg (\log k)^2$ it then is, similar to (8) and (17), routine to see that

$$\mathbb{P}\left(Y_v \geq \delta np \text{ and } \Delta_H \leq (\log k)/(8e_H^2)\right) \leq \left(\frac{e\mu}{\delta np}\right)^{\delta np / \log k} \leq e^{-\delta np / \log k} \ll n^{-1}.$$

Taking a union bound over all vertices v now completes the proof together with estimate (14). \square

It is straightforward, and useful for many applications (see, e.g., [20, 12, 2]), to extend the alteration method to r -uniform hypergraphs, where every edge contains $r \geq 2$ vertices. Indeed, to forbid a given r -uniform hypergraph H , similarly to the graph case ($r = 2$) discussed in Section 1.1, here the idea is to delete edges from a *random r -uniform hypergraph* $G_{n,p}^{(r)}$ (where each of the $\binom{n}{r}$ possible edges appears independently with probability p) to construct an n -vertex r -uniform hypergraph $G \subseteq G_{n,p}^{(r)}$ that is H -free. Defining

$$m_r(H) := \max_{F \subsetneq H} \left(\mathbb{1}_{\{v_F \geq r+1\}} \frac{e_F - 1}{v_F - r} + \mathbb{1}_{\{v_F = r, e_F = 1\}} \frac{1}{r} \right),$$

we say that H is *strictly r -balanced* if $m_r(H) > m_r(F)$ for all $F \subsetneq H$. Noting $G_{n,p} = G_{n,p}^{(2)}$, now the proof of Theorem 1 routinely carries over with only obvious notational changes (where $X_K = |E(G_{n,p}^{(r)}[K])|$ denotes the number of edges of $G_{n,p}^{(r)}$ inside K , and Y_K denotes the number of edges in $E(G_{n,p}^{(r)}[K])$ that are in some H -copy of $G_{n,p}$), yielding the following hypergraphs extension of our main alteration tool Theorem 1.

Theorem 7. *Given $r \geq 2$, let H be a strictly r -balanced r -uniform hypergraph. Then, for any fixed $\delta > 0$, the following holds for all sufficiently large $C \geq C_0 = C_0(\delta, H, r)$ and sufficiently small $0 < c \leq c_0 = c_0(C, \delta, H, r)$. The random r -uniform hypergraph $G_{n,p}^{(r)}$ with $n := \lfloor c(k/\log k)^{m_r(H)} \rfloor$ vertices and edge-probability $p := C(\log k)/k$ whp satisfies $Y_K \leq \delta \binom{k}{r} p$ and $X_K \geq (1 - \delta) \binom{k}{r} p$ for all k -vertex subsets K of $G_{n,p}^{(r)}$.*

Finally, numerous applications [20, 21, 12, 2] of the alteration method require forbidding a finite collection of hypergraphs $\mathcal{H} = \{H_1, \dots, H_s\}$. The crux is that the bounds on Y_K and X_K from Theorem 7 trivially remain valid for $n \leq \lfloor c(k^{r-1}/\log k)^{m_r(H)} \rfloor$. So, applying this result to all forbidden $H_i \in \mathcal{H}$ simultaneously, one can easily obtain a variant of Theorem 7 where Y_K denotes the number of edges in $G_{n,p}^{(r)}[K]$ that are in at least one H_i -copy of $G_{n,p}^{(r)}$ for some $H_i \in \mathcal{H}$; we leave the routine details to the interested reader.

Acknowledgements. We would like to thank Jacob Fox for helpful clarifications regarding [6]. We also thank the referees for useful comments concerning the presentation.

References

- [1] J. Beck. Achievement games and the probabilistic method. In *Combinatorics, Paul Erdős is eighty*, Bolyai Soc. Math. Stud **1** (1993), pp. 51–78.
- [2] T. Bohman, A. Frieze, and D. Mubayi. Coloring H -free hypergraphs. *Rand. Struct. Algor.* **36** (2010), 11–25.
- [3] T. Bohman and P. Keevash. The early evolution of the H -free process. *Invent. Math.* **181** (2010), 291–336.
- [4] T. Bohman and P. Keevash. Dynamic concentration of the triangle-free process. *Rand. Struct. Algor.* **58** (2021), 221–293.
- [5] D. Conlon. On-line Ramsey numbers. *SIAM J. Discrete Math.* **23** (2009), 1954–1963.
- [6] D. Conlon, J. Fox, A. Grinshpun, and X. He. Online Ramsey numbers and the subgraph query problem. In *Building Bridges II*, Bolyai Soc. Math. Stud **28** (2019), pp. 159–194.
- [7] P. Erdős. Graph theory and probability. II. *Canad. J. Math.* **13** (1961), 346–352.
- [8] P. Erdős, R. Faudree, J. Pach, and J. Spencer. How to make a graph bipartite. *J. Combin. Theory Ser. B* **45** (1988), 86–98.
- [9] P. Erdős and P. Tetali. Representations of integers as the sum of k terms. *Rand. Struct. Algor.* **1** (1990), 245–261.
- [10] G. Fiz Pontiveros, S. Griffiths, and R. Morris. The triangle-free process and $R(3, k)$. *Memoirs of the Am. Math. Soc.* **263** (2020), no. 1274.
- [11] J. Fox, X. He, and Y. Wigderson. Ramsey, Paper, Scissors. *Rand. Struct. Algor.* **57** (2020), 1157–1173.
- [12] A. Frieze and D. Mubayi. On the Chromatic Number of Simple Triangle-Free Triple Systems. *Electron. J. Combin.* **15** (2008), Research Paper 121, 27 pp.
- [13] H. Guo and L. Warnke. Packing nearly optimal Ramsey $R(3, t)$ graphs. *Combinatorica* **40** (2020), 63–103.
- [14] H. Guo and L. Warnke. On the power of random greedy algorithms. *European J. Combin.* **105** (2022), 103551.
- [15] S. Janson, T. Łuczak, and A. Ruciński. *Random graphs*. Wiley-Interscience (2000).
- [16] S. Janson and A. Ruciński. The infamous upper tail. *Rand. Struct. Algor.* **20** (2002), 317–342.
- [17] H.A. Kierstead and G. Konjevod. Coloring number and on-line Ramsey theory for graphs and hypergraphs. *Combinatorica* **29** (2009), 49–64.
- [18] J.H. Kim. The Ramsey number $R(3, t)$ has order of magnitude $t^2/\log t$. *Rand. Struct. Algor.* **7** (1995), 173–207.

- [19] M. Krivelevich. Bounding Ramsey numbers through large deviation inequalities. *Rand. Struct. Algor.* **7** (1995), 145–155.
- [20] M. Krivelevich. Approximate set covering in uniform hypergraphs. *J. Algor.* **25** (1997), 118–143.
- [21] M. Krivelevich. On the minimal number of edges in color-critical graphs. *Combinatorica* **17** (1997), 401–426.
- [22] A. Kurek and A. Ruciński. Two variants of the size Ramsey number. *Discuss. Math. Graph Theory* **25** (2005), 141–149.
- [23] M. Kwan, S. Letzter, B. Sudakov, T. Tran. Dense induced bipartite subgraphs in triangle-free graphs. *Combinatorica* **40** (2020), 283–305.
- [24] D. Mubayi and J. Verstraete. A note on pseudorandom Ramsey graphs. *J. Eur. Math. Soc.*, to appear. [arXiv:1909.01461](https://arxiv.org/abs/1909.01461)
- [25] G. Owen. *Game Theory*. 3rd ed., Academic Press, San Diego (1995).
- [26] M. Šileikis and L. Warnke. A counterexample to the DeMarco-Kahn upper tail conjecture. *Rand. Struct. Algor.* **55** (2019), 775–794.
- [27] M. Šileikis and L. Warnke. Upper tail bounds for stars. *Electron. J. Combin.* **27** (2020), Paper no. 1.67, 23 pp.
- [28] J. Spencer. Ramsey’s theorem – a new lower bound. *J. Combin. Theory Ser. A* **18** (1975), 108–115.
- [29] J. Spencer. Asymptotic lower bounds for Ramsey functions. *Discrete Math.* **20** (1977), 69–76.
- [30] B. Sudakov. Ramsey numbers and the size of graphs. *SIAM J. Discrete Math.* **21** (2007), 980–986.
- [31] R. Spöhel, A. Steger, and L. Warnke. General deletion lemmas via the Harris inequality. *J. Combin.* **4** (2013), 251–271.
- [32] L. Warnke. On the missing log in upper tail estimates. *J. Combin. Theory Ser. B* **140** (2020), 98–146.
- [33] L. Warnke. Upper tails for arithmetic progressions in random subsets. *Israel J. Math.* **221** (2017), 317–365.

A Appendix: Lower bound on the upper tail of $|\mathcal{H}_K|$

Given a fixed graph H with $v_H \geq 3$, let us consider a binomial random graph $G_{n,p}$ with edge-probability $p = \Theta((\log k)/k)$ as $k \rightarrow \infty$. Fix a k -vertex subset K of $G_{n,p}$ (which tacitly requires $k \leq n$), and let \mathcal{H}_K denote the collection of all H -copies that have at least one edge inside K . Given $\delta > 0$, we fix v_H disjoint vertex subsets of K , each of size $t := \lceil (\delta \binom{k}{2} p)^{1/v_H} \rceil$. Then $G_{n,p}$ contains with probability $p \binom{v_H}{2} t^2$ a complete v_H -partite subgraph on these v_H sets, which enforces $|\mathcal{H}_K| \geq t^{v_H} \geq \delta \binom{k}{2} p$. It readily follows that

$$\mathbb{P}\left(|\mathcal{H}_K| \geq \delta \binom{k}{2} p\right) \geq p \binom{v_H}{2} t^2 \geq e^{-o(k)},$$

as claimed in Section 1.1 (since $t^2 \cdot \log(1/p) \leq k^{2/v_H + o(1)} \cdot O(\log k) = o(k)$ as $k \rightarrow \infty$).